



A Comprehensive Solution to WAN Performance Issues

A Peribit White Paper

November, 2003

Peribit Networks, Inc.
2855 Bowers Avenue
Santa Clara, CA 95051
USA

+1.866.737.4248
+1.408.330.5600
www.peribit.com
info@peribit.com

Introduction

In an ongoing drive for higher productivity and better customer service, enterprises are rolling out increasingly sophisticated applications such as ERP, CRM, videoconferencing, and voice over IP. As they deploy these applications across their wide area network (WAN) links, enterprise IT managers find that their WAN link capacity and performance make it difficult or impossible to deliver acceptable application performance. Remote facilities may be connected by small WAN links, making it impossible to achieve company-wide rollouts of many new applications. And even when sites are connected with T1 or T3 lines, these links quickly reach capacity as new applications and users are added.

Enterprises typically look to upgrade WAN links to accommodate application and user traffic growth, but bandwidth isn't free and even low-speed international WAN links can cost thousands of dollars per month. Since WAN services are already one of the three top IT expense categories for many global enterprises, managers are understandably reluctant to make matters worse. Moreover, it's simply not possible to upgrade links in some remote locations.

Peribit Networks offers a complete solution to WAN performance problems that doesn't require WAN link upgrades. Peribit Sequence Reducers integrate quickly into existing data center architectures (within 15 minutes in most cases), instantly increase the capacity of existing WAN links up to ten times, and return their costs within months.

Application Performance Constraints

When applications cross the LAN/WAN boundary, they are immediately exposed to several important performance issues. Furthermore, application performance over the WAN is impacted by a variety of factors that must be addressed simultaneously. Solving one limitation improves performance but also uncovers another constraint. Therefore the network manager must take a holistic approach that addresses the three fundamental limitations of WAN application performance: lack of adequate WAN bandwidth, application contention for limited WAN resources, and WAN latency. Without a complete approach, network managers may find that increasing capacity will eliminate a congestion issue and improve performance but it may also reveal that latency is preventing the applications from effectively utilizing this new capacity. Therefore a bandwidth optimization solution must address all facets of application performance over the WAN.

Bandwidth is the capacity of the link. Capacity is the first order problem that impacts application performance. Links that can support standard e-mail and Web browsing traffic often become overloaded when companies add more sophisticated applications.

Contention is the overuse of available capacity by one or more applications. Since most applications will use all of the bandwidth available to them, application performance often suffers on links that support a number of robust applications.

Latency is a delay in moving traffic from one end of a WAN link to the other, in large part caused by router hops and distance. WAN link latency causes delays in setting up connections and providing acknowledgements which slow TCP flows and therefore application performance.




Figure 1. WAN congestion issues: inadequate capacity, application contention, and link latency.

Solution Requirements

To consistently and effectively improve application performance across wide area links, WAN optimization solutions must directly address capacity, contention, and latency. However, in order to be a solution that can be effectively deployed in the enterprise network, there are additional requirements as follows:

- **Simplicity** – The solution must be easy to install and manage. A centralized management station is critical to simplify larger deployments. A solution that avoids an expensive network upgrade but requires significant installation time and ongoing management will not yield effective savings for the organization.
- **Scalability** – The solution must scale to the required number of end points (WAN links) for the enterprise as well as to the necessary link speeds. Frequently this requires the deployment of hundreds of devices that may operate at speeds up to 45 Mbps and beyond.
- **Transparency** – The solution must be transparent to existing network and data center facilities, requiring no reconfiguration of routers, servers, or other devices. The solution must honor existing QoS settings as well as transparently interoperate with a variety of routing protocols and LAN/WAN configurations.
- **Monitoring and Reporting** – Given the products position at the edge of the WAN, these devices should provide visibility into traffic flows with traffic breakdown, top talker, and performance analysis reports. The devices should provide packet capture to help troubleshoot network issues.

Peribit's Approach to WAN Performance Optimization

Peribit's mission is to maximize enterprise application performance by addressing the three key impediments to application performance: congestion, contention and latency. Peribit Sequence Reducers instantly increase WAN capacity, provide essential quality of service features to eliminate contention, and reduce the impact of latency to accelerate application performance. Through a combination of three powerful technologies – Molecular Sequence Reduction (MSR), Essential Quality of Service (Essential QoS), and Packet Flow Acceleration (PFA) – Peribit's Sequence Reducers are the only products that simultaneously address every fundamental WAN performance challenge.

- **MSR Eliminates WAN Congestion:** The first order problem causing application performance issues is limited WAN capacity. MSR reduces the amount of data destined for the WAN link by identifying and eliminating repetitive information. Therefore the existing WAN link can now carry up to 10 times the number of users, sessions or applications without an expensive or time-consuming network upgrade.

- **Essential QoS Prioritizes Applications.** Essential QoS provides predictable application performance by enabling users to allocate this new-found WAN capacity. Peribit will prioritize traffic and allocate bandwidth appropriately among applications to ensure applications perform optimally. Peribit will also honor existing QoS settings in the network preserving this investment.
- **PFA Overcomes Latency:** PFA reduces the impact of latency on enterprise applications by slashing the amount of idle time transmitting or receiving systems wait for delayed packet acknowledgements. PFA increases the responsiveness and throughput of latency-constrained applications up to four times.

The following sections detail these key technologies.

Molecular Sequence Reduction

Peribit's Molecular Sequence Reduction is based on deciphering and understanding the underlying language of the network traffic. MSR then utilizes this knowledge of the network's language to "translate" the original data into a more efficient representation. A second MSR device on the other end of the network then reverse-translates the encoded traffic to generate an identical copy of the original traffic. MSR's advanced language analysis algorithms detect repeated patterns in the traffic and construct a dynamic vocabulary of the key patterns transmitted across the organization. At the same time, MSR learns the dependencies between these patterns to construct a complete knowledge-base of how these basic patterns can be arranged to construct other larger and more complex patterns. This knowledgebase is continuously updated as more patterns and dependencies are learned from the network traffic. At the same time, older patterns that are infrequently used and do not contribute to the data reduction are aged out and replaced with new and more relevant patterns. In this manner, MSR's knowledge-base is kept continuously synchronized with the dynamic language and traffic-characteristics of the network.




Figure 2. MSR identifies and removes repeating patterns

MSR's unique language-learning capabilities, allow this technology to be applied to all types of network traffic at high speed and with very low per-packet latency. The dynamic knowledgebase created by MSR enables the detection and reduction of patterns of any size that are separated by a small or large number of bytes and packets. Hence, both small textual patterns that may be repeated within short distances as well as longer binary patterns that may be separated by hundreds of kilobytes or megabytes of traffic can be detected and reduced by this technology. The unique structure of MSR's knowledge-base also provides very high scalability and allows this technology to operate in real time at very high-speeds.

In contrast to MSR, traditional compression technologies are based on searching and replacing patterns within a fixed historical buffer. This fixed buffer places limits on both the maximum distance between patterns (e.g., only patterns that are less than 2 Kbytes apart will be compressed) as well as the scalability and operating speed of these traditional schemes. As the search buffer size is increased in order to identify more repeated patterns, processing power and latency also increase dramatically. Therefore traditional compression techniques have faced a difficult tradeoff between latency and reduction rates. The result is that while compression is widely available in routers, this tradeoff prevents most enterprises from utilizing this feature across their WAN links. Most networks cannot 'afford' the additional latency (often 30 milliseconds or more) to achieve target reduction levels. By comparison, MSR typically adds less than 2 milliseconds of latency while achieving much greater reduction rates. By employing a completely new approach to dynamically learn and intelligently store a knowledge base of repeated patterns, MSR overcomes these limitations to provide significantly greater pattern detection and reduction with virtually unlimited scalability.

Essential QoS

Peribit's Essential Quality of Service (QoS) features give network managers an easy and powerful way to eliminate application contention and ensure optimum performance. MSR creates more WAN capacity for all applications, thereby breaking the "winner" and "loser" limitation of traditional QoS solutions where one application is prioritized over another. (Peribit improves the performance of all applications, so there are no losers.)

Peribit's essential QoS features allow network managers to prioritize traffic and allocate application bandwidth. In networks where QoS has already been deployed, Peribit's Sequence Reducers transparently honor and preserve existing settings, so adding Peribit Sequence Reducers requires no reconfiguration of other QoS equipment.




Figure 3. Application contention without Essential QoS




Figure 4. Application contention eliminated with Essential QoS

Peribit's approach to QoS is "top-down," making it much simpler than other, more complicated "bottom-up" approaches. Traditional approaches required detailed analysis of traffic flows to build up a complete QoS solution. With Peribit's QoS, traffic priorities and bandwidth allocations are defined by both

application and corporate location. This allows the system administrator to simply configure QoS settings using simple drop down menus assigning traffic classes and then allocating percentages of WAN link capacity. The Sequence Reducers then dynamically apply these policies to reduced traffic to ensure that each application receives appropriate bandwidth allocation.

The result is that Peribit's Essential QoS allows network managers to easily prioritize and allocate the additional bandwidth created by MSR. With Peribit, network managers can ensure higher priority applications get the bandwidth they need to perform optimally while allowing lower priority applications to continue to operate as before. With an easy, intuitive interface, Peribit's essential QoS simplifies deployment and does not create an additional management burden.

Packet Flow Acceleration

Packet Flow Acceleration (PFA) dramatically reduces application response time by overcoming the effects of WAN latency. PFA delivers LAN-like application performance over the WAN. PFA dramatically speeds application response time by eliminating delays in TCP/IP packet flows. Typically, a transmitting system sends a small number of packets, the receiving system acknowledges them, and then the transmitting system sends more packets. In this scenario, the transmitting system sends initial packets and then waits for acknowledgment before sending more packets, and the receiving system sends acknowledgments and then waits for the transmitting system to send more packets.



Figure 5. Typical TCP/IP transmission where transfers are slowed waiting for acknowledgements

These wait cycles cause transmission delays, which vary based on the severity of several factors such as natural propagation delays on long distance network links, router hop delays introduced in private networks or VPNs, or delays from very high latency links like satellite connections. In any case, such delays stall the efficient flow of traffic across the WAN.

PFA minimizes these transmission delays by pre-loading subsequent packets to the remote location. Thus, the receiving system does not wait for additional packets, but can continually process them. And, as the receiving system processes its larger payload of packets, its acknowledgments reach the transmitting system, which then sends additional packets while the receiving system is still processing. As a result, PFA reduces the time it takes to complete a WAN transmission by as much as 75 percent, slashing file transfer times and speeding response times for capacity constrained applications by up to 4 times.



Figure 6. TCP/IP transmission with PFA loads more packets, reduces idle time and overcoming impact of latency

Monitoring and Reporting

All Peribit Sequence Reducers include powerful monitoring and reporting features that deliver real-time views of network performance. With Peribit's intuitive graphical interface, users can quickly generate standard reports on "top talkers" or other traffic attributes to gain instant visibility into network traffic and performance characteristics. And, because Peribit's MSR and PFA technologies rely on full packet inspection, users can view packet-level traffic statistics for detailed troubleshooting, or export traffic data to other programs for further analysis and reporting.

Products

Peribit's MSR, essential QoS, and PFA technologies are offered in a family of Sequence Reducers – rack-mountable network devices that interface to the network as a transparent 2-port LAN switch, situated between the aggregation switch and the WAN router. Peribit Sequence Reducers are LAN based, so they can support any physical transmission medium or topology, including fractional T1/E1, T1, T3/E3, DSL, frame relay, ISDN, ATM, or VPN service.

Peribit Networks offers the following products:



SR-20

Supports up to 2 Mbps circuits and manages up to 5 end-points with other Sequence Reducers for remote offices or smaller sites - 10/100 BaseT copper Ethernet interfaces



SR-50

Supports up to 45 Mbps circuits and manages up to 120 end-points with other Sequence Reducers - 10/100 BaseT copper Ethernet interfaces



SR-55

Supports up to 45 Mbps circuits and manages up to 120 end-points with other Sequence Reducers – 10/100/1000 BaseT copper Ethernet interfaces



SR-80

Supports up to 45 Mbps circuits and manages up to 320 end-points with other Sequence Reducers – 10/100/1000 BaseT copper and 1000 BaseSX optical interfaces



Central Management System (CMS)

Software to centrally deploy, manage, and monitor large networks of Sequence Reducers with advanced configuration management and monitoring.

Summary

Only Peribit's Sequence Reducers overcome all three roadblocks to WAN link performance: bandwidth, contention, and latency. Peribit's MSR technology multiplies available bandwidth up to 10 times by identifying and removing redundant data patterns before they traverse the WAN. Peribit's Essential QoS features deliver prioritization and bandwidth allocation capabilities to ensure optimum performance on a per-application basis, while seamlessly integrating with any existing network-based end-to-end QoS solution. Finally, Peribit's Packet Flow Acceleration (PFA) technology further improves application performance and response times by minimizing the delays and inefficiencies caused by latency on TCP-based communications.

Peribit can demonstrate how the combination of MSR, QoS, and PFA can optimize your application performance and help you reduce operational expenditures. Contact Peribit at www.peribit.com or info@peribit.com.

Peribit, Molecular Sequence Reduction, Packet Flow Acceleration, and the Peribit logo are trademarks of Peribit Networks.